REMARKS

The Office Action dated November 21, 2005 has been received and carefully noted. The above amendments to claims 1, 8, 14-17, 20, 22, 25-27, 31, 33, 35, 37, and 40-42 and the following remarks, are submitted as a full and complete response thereto. Support for the amended recitations may be found, at least, on page 62, lines 12-30 of the specification and FIGS. 40-42 and corresponding descriptions. No new matter is being presented, and approval and entry are respectfully requested.

Claims 1-3, 5-12, 14-23, 25-35, 37-38, and 40-42 stand rejected and pending and under consideration.

OBJECTIONS TO THE CLAIMS:

In the Office Action, at page 2, claims 22 and 33 were objected to because "said memory interface," according to the office action, should be changed to "said external memory interface." Applicants respectfully traverse such objection. Claim 22 depends from independent claim 14 and dependent claim 33 depends from independent claim 25. Both independent claims recite "a memory interface." Accordingly, independent claims 22 and 25 provide proper antecedent support for the features recited in claims 22 and 33, respectively. Accordingly, it is respectfully requested that the objections to these claims be withdrawn.

Claims 41 and 42 were objected to for informalities. Claims 41 and 42 have been amended to resolve such informalities and, accordingly, it is respectfully requested that the objections to these claims be withdrawn.

REJECTION UNDER 35 U.S.C. § 102:

In the Office Action, at page 2, claims 1-3, 6-12, 14-23, 25-35, 37-38 and 40-42 were rejected under 35 U.S.C. § 102 as being anticipated by U. S. Patent No. 6,246,680 to Muller et al. ("Muller"). The Office Action took the position that Muller describes all the recitations of independent claims 1, 14, 15, 16, 20, 25, 26, 27, 31, 35, 37, 40, and 41 and related dependent claims. It is respectfully asserted that, for at least the reasons provided herein below, Muller fails to teach or suggest the recitations of the pending claims. Reconsideration is requested.

Independent claim 1, upon which claims 2, 3, and 5-12 are dependent, recites a network switch, the network switch comprising at least one data port interface supporting a plurality of data ports, a submodule adding an interstack tag into data to keep track of a stack count to prevent looping of the data, at least one stack link interface comprising a bi-directional gigabit stack link interface configured to transmit the data between the network switch and other network switches to create a predetermined configuration, and a CPU interface, the CPU interface configured to communicate with a CPU. A memory management unit is in communication with the at least one data port interface and the at

least one stack link interface, and a memory interface is in communication with the at least one data port interface and the at least one stack link interface, wherein the memory interface is configured to communicate with a memory. A communication channel is provided for communicating data and messaging information between the at least one data port interface, the at least one stack link interface, the memory interface, and the memory management unit. The memory management unit is configured to route data received from each of the at least one data port interface and the at least one stack link interface to the memory interface.

Independent claim 14, upon which claims 21-23 are dependent, recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks comprises at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration, and a CPU interface configured to communicate with a CPU. A memory management unit is in communication with the at least one data port interface and the predetermined number of stack link interfaces, a memory interface is in communication with the at least one data port interface and the

predetermined number of stack link interfaces, wherein the memory interface is configured to communicate with a memory, and a communication channel is provided for communicating data and messaging information between the at least one data port interface, the predetermined number of stack link interfaces, the memory interface, and the memory management unit.

Claim 15 recites a scalable network switch that includes a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch buildings blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data and a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data. The at least one of the predetermined number of switch building blocks also includes a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration. Also, the predetermined number of stack link interfaces is configured to be one less than the predetermined number of switch building blocks.

Claim 16 recites a scalable network switch that includes a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data and a submodule

adding an interstack tag into the data to keep track of a stack count to prevent looping of the data. The at least one of the predetermined number of switch building blocks also includes a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration. Also, the at least one data port interface includes at least one first data port interface supporting a plurality of first data ports transmitting and receiving data at a first data rate and at least one second data port interface supporting at least one second data port transmitting and receiving data at a second rate.

Claim 20 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks comprises at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration, wherein each of the predetermined number of stack link interfaces further

comprise a gigabit stack link interface configured to transmit and receive data from another gigabit stack link interface on another switch building block.

Claim 25 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration. At least one of the predetermined number of switch building blocks further includes a CPU interface configured to communicate with a CPU, a memory management unit in communication with the at least one data port interface and the predetermined number of stack link interfaces, a memory interface in communication with the at least one data port interface and the predetermined number of stack link interfaces, wherein the memory interface is configured to communicate with a memory, and a communication channel, the communication channel for communicating data and messaging information between the at least one data port interface, the predetermined number of stack link interfaces, the memory interface, and the memory management unit.

Claim 26 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration, wherein the predetermined number of stack link interfaces is configured to be one less than the predetermined number of switch building blocks.

Claim 27 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks comprises: at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration. The at least one data port interface further comprises: at least one first

data port interface supporting a plurality of first data ports transmitting and receiving data at a first data rate, and at least one second data port interface supporting at least one second data port transmitting and receiving data at a second rate.

Claim 31 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration, wherein each of the predetermined number of stack link interfaces further comprise a gigabit stack link interface configured to transmit and receive data from another gigabit stack link interface on another building block.

Claim 35 recites a scalable network switch, the scalable network switch comprising a predetermined number of switch building blocks interconnected in a meshed configuration. At least one of the predetermined number of switch building blocks includes at least one data port interface supporting a plurality of data ports for transmitting and receiving data, a submodule adding an interstack tag into the data to keep track of a stack count to prevent looping of the data, and a predetermined number of

stack link interfaces comprising bi-directional gigabit stack link interfaces configured to transmit the data between one of the predetermined number of building blocks and another of the predetermined number of building blocks to create a predetermined configuration, the scalable network switch further comprising a physical layer transceiver in connection with at least one of the plurality of data ports.

Claim 37 recites a method of stacking network switches. The method includes providing a plurality of clustered switch blocks, and interconnecting each one of the plurality of clustered switch blocks to another one of the plurality of clustered switch blocks. Interconnection of the plurality of clustered switch blocks forms a stack of clustered switch blocks, wherein the step of providing a plurality of clustered switch blocks further includes providing a predetermined number of switch building blocks, interconnecting each of the predetermined number of switch building blocks to every other one of the predetermined number of switch building blocks in a meshed configuration, and adding an interstack tag into data received to keep track of a stack count to prevent looping of the data. Each of the predetermined number of switch building blocks is interconnected to every other one of the predetermined number of switch blocks through an individual stack link.

Claim 40 recites a method of stacking network switches, the method comprising the steps of: providing a plurality of clustered switch blocks, and interconnecting each one of the plurality of clustered switch blocks to another one of the plurality of clustered switch blocks. Interconnection of the plurality of clustered switch blocks forms a stack

of clustered switch blocks. The receiving step further comprises the steps of: receiving a packet on at least one of a data port interface and a stack link interface, adding an interstack tag into the packet to keep track of a stack count to prevent looping of the packet, and storing the packet in a memory in accordance with a predetermined algorithm by allocating memory locations in an internal memory and in an external memory based upon an amount of internal memory available for an egress port of the clustered network switch from which the packet is to be transmitted.

Claim 41 recites a method of stacking network switches, the method comprising the steps of: providing a plurality of clustered switch blocks, and interconnecting each one of the plurality of clustered switch blocks to another one of the plurality of clustered switch blocks. Interconnection of the plurality of clustered switch blocks forms a stack of clustered switch blocks. The forwarding step further includes receiving a packet, adding an interstack tag into the packet to keep track of a stack count to prevent looping of the packet, determining if the destination address of the packet corresponds to a port in the clustered network switch, forwarding the packet to the port corresponding to the destination address if the destination address is determined to correspond to a port in the clustered network switch, determining if the destination address of the packet corresponds to a port on another clustered network switch across a stack, forwarding the packet to a stack link if the destination address is determined to correspond to a port on the another clustered network switch across the stack, and transmitting the packet across

the stack to the another clustered network switch if the destination address of the packet corresponds to a port on the another clustered network switch across the stack.

As will be discussed below, Muller fails to disclose or suggest the elements of any of the presently pending claims.

Muller generally describes an architecture for an integrated network element building block including a network interface 205 with multiple ports for transmitting and receiving packets over a network. See FIG. 2 and column 4, lines 5-67. A shared memory manager 220 is provided to allocate and deallocate buffers in the packet buffer storage on behalf of the network interface and other clients of the packet buffer storage. The network device 205 building block further includes a switch fabric 210 which is coupled to the network interface 205. The switch fabric 210 provides forwarding decisions for received packets. A given forwarding decision includes a list of ports upon which a particular received packet is to be forwarded.

In Muller, input packet processing includes the following: (1) receiving and verifying incoming Ethernet packets, (2) modifying packet headers when appropriate, (3) requesting buffer pointers from the shared memory manager 220 for storage of incoming packets, (4) requesting forwarding decisions from the switch fabric block 210, (5) transferring the incoming packet data to the shared memory manager 220 for temporary storage in an external shared memory 230, and (5) upon receipt of a forwarding decision, forwarding the buffer pointer(s) to the output port(s) indicated by the forwarding decision. See column 4, lines 43-60. Output packet processing may be performed by one

or more output ports of the network interface 205. Output processing includes requesting packet data from the shared memory manager 220, transmitting packets onto the network, and requesting deallocation of buffer(s) after packets have been transmitted.

However, Muller fails to teach or suggest, at least, "a submodule adding an interstack tag into data to keep track of a stack count to prevent looping of the data," as recited in independent claims 1, 14-16, 20, 25-27, 31, and 35. Similarly, Muller is silent as to teaching or suggesting, at least, "adding an interstack tag into data received to keep track of a stack count to prevent looping of the data," as recited in independent claims 37, 40, and 41. Muller does not recognize that, in order to prevent looping of the data or the packet being received and processed, an interstack tag may be added in the data or the packet. Muller simply provides that the input packet processing includes modifying packet headers when appropriate, without providing any description regarding as to when it is appropriate to modify the input packet. Also, Muller does not provide a description regarding as to what is meant by "modifying" the packet header.

Accordingly, in view of the foregoing, it is respectfully requested that independent claims 1, 14-16, 20, 25-27, 31, 35, 37, 40, and 41 and related dependent claims be allowed.

REJECTION UNDER 35 U.S.C. § 103:

In the Office Action, at page 6, claims 5, 23, and 34 were rejected under 35 U.S.C. § 103 as being unpatentable over Muller and U.S. Patent No. 6,775,290 to Merchant et al. ("Merchant"). The Office Action took the position that Muller and Merchant discloses all the aspects of dependent claims 5, 23, and 34. The rejection is traversed and reconsideration is requested.

As will be discussed below, Muller and Merchant fail to disclose or suggest the elements of any of the presently pending claims.

Dependent claim 5 depends from independent claim 1 and recites the additional features of "a variable sized address resolution logic table; and a variable sized VLAN table, wherein said variable sized address resolution logic table and said variable sized VLAN table is in communication with said memory management unit, said at least one stack link interface, and said at least one data port interface." Because the combination of Muller and Merchant must teach, individually or combined, all the recitations of the base claim and any intervening claims of dependent claim 5, the arguments presented above supporting the patentability of independent claim 1 over Muller are incorporated herein.

Merchant generally describes a method to enable a port of a network switch to support connections with multiple VLANs. See column 1, lines 50-55. Each multiport switch 12 includes a media access control (MAC) module 20 that transmits and receives

data packets to and from 10/100 Mb/s physical layer (PHY) transceivers 16 via respective reduced media independent interfaces (RMII) 18 according to IEEE 802.3u protocol. Each multiport switch 12 also includes a gigabit MAC 24 for sending and receiving data packets to and from a gigabit PHY 26 for transmission to the gigabit node 22 via a high speed network medium 28. See column 3, lines 38-47.

However, Merchant does not cure the deficiencies of Muller. Similarly to Muller, Merchant fails to teach or suggest, at least, "a submodule adding an interstack tag into data to keep track of a stack count to prevent looping of the data," as recited in independent claim 1. Instead, Merchant describes, at most, a frame that may include a VLAN tag header that identifies the frame as information destined to one or more members of a prescribed group of stations. See column 5, lines 45-48. A combination of Muller and Merchant would not provide for the entire claim recitations of independent claim 1, and, accordingly, dependent claim 5.

Accordingly, in view of the foregoing, it is respectfully requested that independent claim 1 and related dependent claim 5 be allowed.

CONCLUSION:

In view of the above, Applicant respectfully submits that the claimed invention recites subject matter which is neither disclosed nor suggested in the cited prior art. Applicant further submits that the subject matter is more than sufficient to render the claimed invention unobvious to a person of skill in the art. Applicant therefore

respectfully requests that each of claims 1-3, 5-12, 14-23, 25-35, 37-38, and 40-42 be found allowable and this application passed to issue.

If for any reason the Examiner determines that the application is not now in condition for allowance, it is respectfully requested that the Examiner contact, by telephone, the applicant's undersigned attorney at the indicated telephone number to arrange for an interview to expedite the disposition of this application.

In the event this paper is not being timely filed, the Applicant respectfully petitions for an appropriate extension of time.

In the event this paper is not being timely filed, the applicants respectfully petition for an appropriate extension of time. Any fees for such an extension together with any additional fees may be charged to Counsel's Deposit Account 50-2222.

Respectfully submitted,

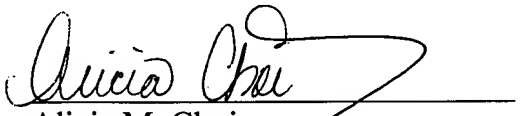

Alicia M. Choi
Registration No. 46,621

Customer No. 32294
SQUIRE, SANDERS & DEMPSEY LLP
14TH Floor
8000 Towers Crescent Drive
Tysons Corner, Virginia 22182-2700
Telephone: 703-720-7800
Fax: 703-720-7802

AMC:jkm